

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	0	generat\$4 same customer\$1 same record\$1 same table\$1 same preferenc\$4 same deliver\$4 same output\$1	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/07 14:38
L2	1	access\$4 same content\$1 same file\$1 same quer\$4 same (database\$1 or db\$1 or (data adj base\$1)) same customer\$1 same receiv\$4 same record\$1	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/07 14:41
L3	0	transmi\$5 same output\$1 same determin\$4 same deliver\$4 same option\$1 same customer\$1 same record\$1	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/07 14:42

09/38, 267



US Patent & Trademark Office

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

access\$4 same content\$1 same file\$1 same quer\$4 same (dat

SEARCH

[Feedback](#) [Report a problem](#) [Sati](#)

Terms used

access\$4 same content\$1 same file\$1 same quer\$4 same database\$1 or db\$1 or data adj base\$1 same customer\$1 sam

Sort results by

Display results

[Save results to a Binder](#)

[Search Tips](#)

☐ [Open results in a new window](#)

Try an [Advanced Sea](#)

Try this search in [The](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

1 [Comparison of access methods for time-evolving data](#)

Betty Salzberg, Vassilis J. Tsotras

June 1999

ACM Computing Surveys (CSUR), Volume 31 Issue 2

Full text available: pdf(529.53 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper compares different indexing techniques proposed for supporting efficient access to temporal data. collection of important performance criteria, including the space consumed, update processing, and query tim The comparison is based on worst-case analysis, hence no assumptions on data distribution or query frequenc number of methods have the same asymptotic worst-case behavior, features in the methods tha ...

Keywords: I/O performance, access methods, structures, temporal databases

2 [Modeling and performance evaluation of physical data base structures](#)

S. B. Yao

October 1976

Proceedings of the annual conference

Full text available: pdf(543.60 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A generalized file organization model and performance evaluation system is developed for estimating the perf base structures. Performance evaluation results based on the cost functions of the model and system are pres measures are compared with the results of previous simulation models for indexed sequential, multi-list, and analytic approach makes the costs of evaluation very low. Consequently, many evaluations m ...

3 [Data broadcasting strategies over multiple unreliable wireless channels](#)

Hong V. Leong, Antonio Si

December 1995

Proceedings of the fourth international conference on Information and knowledge man

Full text available: pdf(1.07 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

4 [A prototype implementation of the SQL Ada module extension \(SAME\) method](#)

Allison LeClair, Susan Phillips

December 1990

Proceedings of the conference on TRI-ADA '90

Full text available: pdf(1.20 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)


As Ada becomes more widespread, the ability to access commercial database technologies through Ada system issue. Researchers throughout our industry are investigating interface approaches between Ada and these tec bindings between Ada and SQL, a relational data base language. This paper presents a recent implementation SQL Ada Module Extension (SAME) method.

5 Object orientation in multidatabase systems

Evaggelia Pitoura, Omran Bukhres, Ahmed Elmagarmid

June 1995

ACM Computing Surveys (CSUR), Volume 27 Issue 2

Full text available:  pdf(4.85 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#)

A multidatabase system (MDBS) is a confederation of preexisting distributed, heterogeneous, and autonomous. A recent proliferation of research suggesting the application of object-oriented techniques to facilitate designing and implementing MDBSs. Although this approach seems promising, the lack of a general framework development. The goal of this paper is to provide a concrete analysis and categorization of the various ...

Keywords: distributed objects, federated databases, integration, multidatabases, views

6 Compiler-based I/O prefetching for out-of-core applications

Angela Demke Brown, Todd C. Mowry, Orran Krieger

May 2001

ACM Transactions on Computer Systems (TOCS), Volume 19 Issue 2

Full text available:  pdf(499.03 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#)

Current operating systems offer poor performance when a numeric application's working set does not fit in memory. Programmers who wish to solve "out-of-core" problems efficiently are typically faced with the onerous task of using explicit I/O operations (e.g., read/write). In this paper, we propose and evaluate a fully automatic technique. A programmer from this task, provides high performance, and requires only minimal ...

Keywords: compiler optimization, prefetching, virtual memory

7 On the analytical modeling of database concurrency control

Philip S. Yu, Daniel M. Dias, Stephen S. Lavenberg

September 1993 **Journal of the ACM (JACM)**, Volume 40 Issue 4

Full text available:  pdf(2.75 MB)


Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#)

The Concurrency Control (CC) scheme employed can profoundly affect the performance of transaction-processing. A simple unified approximate analysis methodology to model the effect on system performance of data contention schemes and for different system structures is developed. This paper concentrates on modeling data contention. In other papers, the solutions of the data contention model are coupled with a standard hard ...

8 Consistency and orderability: semantics-based correctness criteria for databases

Divyakant Agrawal, Amr El Abbadi, Ambuj K. Singh

September 1993 **ACM Transactions on Database Systems (TODS)**, Volume 18 Issue 3

Full text available:  pdf(1.92 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#)

The semantics of objects and transactions in database systems are investigated. User-defined predicates called used to specify user programs. Three new correctness criteria are proposed. The first correctness criterion considers the users' specifications and admits nonserializable executions that are acceptable to the users. Integrity constraints maintained through consistency assertions. The ...

Keywords: concurrency control, object-oriented databases, semantics, serializability theory

9 Anatomy of a native XML base management system

T. Fiebig, S. Helmer, C.-C. Kanne, G. Moerkotte, J. Neumann, R. Schiele, T. Westmann

December 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue

Full text available:  pdf(300.97 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

Several alternatives to manage large XML document collections exist, ranging from file systems over relational to specifically tailored XML base management systems. In this paper we give a tour of Natix, a database management system from scratch for storing and processing XML data. Contrary to the common belief that management of XML data for traditional databases like relational systems, we illustrate how almost every component in a ...

Keywords: Database, XML

10 Database Reorganization—Principles and Practice

Gary H. Sockut, Robert P. Goldberg

December 1979 **ACM Computing Surveys (CSUR)**, Volume 11 Issue 4

Full text available:  pdf(1.89 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

11 Deductive database languages: problems and solutions

Mengchi Liu

March 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 1

Full text available:  pdf(254.50 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Deductive databases result from the integration of relational database and logic programming techniques. They remain inherent in this simple synthesis from the language point of view. In this paper, we discuss these problems and aspects: complex values, object orientation, higher-orderness, and updates. In each case, we examine four types of the corresponding issues.

Keywords: complex object databases, deductive databases, inheritance, logic programming, nested relational databases

12 Pipeline Architecture

C. V. Ramamoorthy, H. F. Li

January 1977 **ACM Computing Surveys (CSUR)**, Volume 9 Issue 1

Full text available:  pdf(3.53 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

13 Experiences with the Amoeba distributed operating system

Andrew S. Tanenbaum, Robbert van Renesse, Hans van Staveren, Gregory J. Sharp, Sape J. Mullender

December 1990 **Communications of the ACM**, Volume 33 Issue 12

Full text available:  pdf(2.71 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Amoeba project is a research effort aimed at understanding how to connect multiple computers in a seamless way [31]. The basic idea is to provide the users with the illusion of a single powerful timesharing system, when, in fact, the system is implemented on a collection of machines, potentially distributed among several countries. This research has led to the implementation of the Amoeba distributed operating system, which is being used as a prototype and vehicle for further research.

14 Research sessions: implementation techniques: Fractal prefetching B⁺-Trees: optimizing both cache and I/O

Shimin Chen, Phillip B. Gibbons, Todd C. Mowry, Gary Valentin

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available:  pdf(1.49 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

B⁺-Trees have been traditionally optimized for I/O performance with disk pages as tree nodes. Recently, researchers have developed types of B⁺-Trees optimized for CPU cache performance in main memory environments, where the tree node size is small. Unfortunately, due primarily to this large discrepancy in optimal node sizes, existing disk-optimized B⁺-Trees exhibit poor performance while cache-optimized B⁺-Trees exhibit good performance.

15 Discovering roll-up dependencies

Jef Wijsen, Raymond T. Ng, Toon Calders

August 1999 **Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery in data mining**

Full text available:  pdf(994.24 KB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

16

Intelligent database caching through the use of page-answers and page-traces

Nabil Kamel, Roger King

December 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 4

Full text available:  pdf(3.08 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper a new method to improve the utilization of main memory systems is presented. The new method main memory a number of query answers, each evaluated out of a single memory page. To this end, the idea traces are formally described and their properties analyzed. The query model used here allows for selection, p queries as well as arbitrary combinations. We also show how to apply the approach under update ...

Keywords: artificial intelligence, databases, page access

17 An analytical model of the working-set sizes in decision-support systems

Magnus Karlsson, Per Stenström

June 2000

ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2000 ACM SIGM conference on Measurement and modeling of computer systems, Volume 28 Issue 1

Full text available:  pdf(997.83 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents an analytical model to study how working sets scale with database size and other applica support systems (DSS). The model uses application parameters, that are measured on down-scaled database miss ratios for executions of large databases. By applying the model to two database engines and typical DSS large databases, the most performance-critical working set is small and is ...

18 Concurrency control: methods, performance, and analysis

Alexander Thomasian

March 1998

ACM Computing Surveys (CSUR), Volume 30 Issue 1

Full text available:  pdf(427.18 KB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: Markov chains, adaptive methods, concurrency control, data contention, deadlocks, flow diagram concurrency control, queueing network models, restart-oriented locking methods, serializability, thrashing, t processing, wait depth limited methods

19 Parallel changes in large-scale software development: an observational case study

Dewayne E. Perry, Harvey P. Siy, Lawrence G. Votta

July 2001

ACM Transactions on Software Engineering and Methodology (TOSEM), Volume 10 Issue 3

Full text available:  pdf(361.44 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

An essential characteristic of large-scale software development is parallel development by teams of developer development is structured and supported has a profound effect on both the quality and timeliness of the prod observational case study in which we collect and analyze the change and configuration management history o the boundaries of, and to understand the nature of, the problems encountered in parallel development. ...

Keywords: change management, merging interfering and noninterfering versions, parallel versions, parallel/ integration

20 TIMBER: A native XML database

H. V. Jagadish, S. Al-Khalifa, A. Chapman, L. V. S. Lakshmanan, A. Nierman, S. Paparizos, J. M. Patel, D. Srivas Wu, C. Yu

December 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue

Full text available:  pdf(268.39 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

This paper describes the overall design and architecture of the Timber XML database system currently being i of Michigan. The system is based upon a bulk algebra for manipulating trees, and natively stores XML. New ac developed to evaluate queries in the XML context, and new cost estimation and query optimization techniques We present performance numbers to support some of our design decisions. We believe that the key in ...

Keywords: Algebra, Document management, Hierarchical, Query processing, Semi-structured

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, In
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Play](#)



US Patent & Trademark Office

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

access\$4 same content\$1 same file\$1 same quer\$4 same (dat

SEARCH

[Feedback](#) [Report a problem](#) [Sati](#)

Terms used

access\$4 same content\$1 same file\$1 same quer\$4 same database\$1 or db\$1 or data adj base\$1 same customer\$1 sam

Sort results by

Display results

[Save results to a Binder](#)

[Search Tips](#)

☐ [Open results in a new window](#)

[Try an Advanced Sea](#)

[Try this search in The](#)

Results 21 - 40 of 200

Best 200 shown

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

21 [Evolution of Data-Base Management Systems](#)

James P. Fry, Edgar H. Sibley

January 1976 **ACM Computing Surveys (CSUR)**, Volume 8 Issue 1

Full text available: pdf(2.63 MB)

Additional Information: [full citation](#), [references](#), [citing](#), [index terms](#)

22 [Long-duration transaction support in design databases](#)

Waldemar Wiczerzycki

December 1995 **Proceedings of the fourth international conference on Information and knowledge man**

Full text available: pdf(852.24 KB)

Additional Information: [full citation](#), [references](#), [index terms](#)

23 [Universal service-providers for database private information retrieval \(extended abstract\)](#)

Giovanni Di-Crescenzo, Yuval Ishai, Rafail Ostrovsky

June 1998 **Proceedings of the seventeenth annual ACM symposium on Principles of distributed comp**

Full text available: pdf(1.70 MB)

Additional Information: [full citation](#), [references](#), [citing](#), [index terms](#)

24 [A database model for object dynamics](#)

M. P. Papazoglou, B. J. Krämer

May 1997 **The VLDB Journal – The International Journal on Very Large Data Bases**, Volume 6 Issue 2

Full text available: pdf(313.64 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

To effectively model complex applications in which constantly changing situations can be represented, a datab support the runtime specification of structural and behavioral nuances for objects on an individual or group ba role mechanism as an extension of object-oriented databases to support unanticipated behavioral oscillations many types and share a single object identity. A role refers to the ability to repres ...

Keywords: Dynamic class hierarchy, Dynamic object re-classification, Object migration, Object role model, O systems

25 [Special issue of the lexicon: Tools and methods for computational lexicology](#)

Roy J. Byrd, Nicoletta Calzolari, Martin S. Chodorow, Judith L. Klavans, Mary S. Neff, Omneya A. Rizk

July 1987 **Computational Linguistics**, Volume 13 Issue 3-4

Full text available:

Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#)

 pdf(2.49 MB)  [Publisher Site](#)

This paper presents a set of tools and methods for acquiring, manipulating, and analyzing machine-readable d detailed examples of the use of these tools and methods for particular analyses. A novel aspect of our work is processing of multiple machine-readable dictionaries. Our examples describe analyses of data from Webster's Dictionary, the Longman Dictionary of Contemporary English, the Collins bilingual dictionaries, t ...

26 A distributed object-oriented database system supporting shared and private databases

Won Kim, Nat Ballou, Jorge F. Garza, Darrell Woelk

January 1991 **ACM Transactions on Information Systems (TOIS)**, Volume 9 Issue 1

Full text available:  pdf(1.58 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#) ____

ORION-2 is a commercially available, federated, object-oriented database management system designed and major architectural innovation in ORION-2 is the coexistence of a shared database and a number of private da database is accessible to all authorized users of the system, while each private database is accessible to only distributed database system with a shared database and private databases for individual users is a natu ...

Keywords: client-server architecture, federated databases, object-oriented databases

27 Some mathematical aspects on syntactic discription

Itiroo Sakai

May 1965

Proceedings of the 1965 conference on Computational linguistics

Full text available:  pdf(2.32 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#)

The purpose of this paper is to help linguists contract a consistent, sufficient and less redundant syntax of lan corresponds to an expression or an utterance: it may be a natural text, a string of morphemes, a tree structu representation. A sharp distinction is made between the syntactic function which is an attribute of string and a set of strings. Syntactic function of a continuous or discontinuous string is defined as ...

28 Multikey access methods based on superimposed coding techniques

R. Sacks-Davis, A. Kent, K. Ramamohanarao

November 1987 **ACM Transactions on Database Systems (TODS)**, Volume 12 Issue 4

Full text available:  pdf(3.71 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#) ____

Both single-level and two-level indexed descriptor schemes for multikey retrieval are presented and compared using superimposed coding techniques and stored using a bit-inversion technique. A fast-batch insertion algor forming the bit-inverted file is less than one disk access per record is presented. For large data files, it is show implementation is generally more efficient for queries with a small number of matchin ...

29 Greed sort: optimal deterministic sorting on parallel disks

Mark H. Nodine, Jeffrey Scott Vitter

July 1995 **Journal of the ACM (JACM)**, Volume 42 Issue 4

Full text available:  pdf(1.00 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#)

We present an algorithm for sorting efficiently with parallel two-level memories. Our main result is an elegant deterministic algorithm for external sorting with D disk drives. This result answers in the affirmative the open Shriver of whether an optimal algorithm exists that is deterministic. Our measure of performance is the numb (I/O) operations, in which each of the

Keywords: I/O complexity, merge sort, parallel I/O, parallel disks

30 Locking Primitives in a Database System

Henry F. Korth

January 1983 **Journal of the ACM (JACM)**, Volume 30 Issue 1

Full text available:  pdf(1.61 MB)

Additional Information: [full citation](#), [references](#), [citings](#), [index terms](#)

31 Scalable packet classification

Florin Baboescu, George Varghese

August 2001

ACM SIGCOMM Computer Communication Review , Proceedings of the 2001 conference technologies, architectures, and protocols for computer communications, Volume 31 Issue


Full text available:  pdf(242.61 KB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

32 Toward memory-based reasoning

Craig Stanfill, David Waltz

December 1986 **Communications of the ACM**, Volume 29 Issue 12

Full text available:  pdf(1.66 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The intensive use of memory to recall specific episodes from the past—rather than rules—should be the found

33 Parity logging disk arrays

Daniel Stodolsky, Mark Holland, William V. Courtright, Garth A. Gibson

August 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 3

Full text available:  pdf(1.98 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Parity-encoded redundant disk arrays provide highly reliable, cost-effective secondary storage with high performance writes. Their performance on small writes, however, is much worse than mirrored disks—the traditional, high organization for secondary storage. Unfortunately, small writes are a substantial portion of the I/O workload of demanding applications such as on-line transaction processing. This paper presents

Keywords: RAID, disk arrays

34 The use of normal multiplication tables for information storage and retrieval

Dalia Motzkin

March 1979 **Communications of the ACM**, Volume 22 Issue 3

Full text available:  pdf(1.38 MB)

Additional Information: [full citation](#), [abstract](#), [references](#)

This paper describes a method for the organization and retrieval of attribute based information systems, using a table as a directory for the information system. Algorithms for the organization and retrieval of information are particularly suitable for queries requesting a group of information items, all of which possess a particular set of some other attributes as well). Several examples are given; the results with ...

Keywords: information retrieval, inverted files, multiattribute retrieval, multilist file, normal multiplication table, space economy

35 Support for repetitive transactions and ad hoc queries in System R

D. D. Chamberlin, M. M. Astrahan, W. F. King, R. A. Lorie, J. W. Mehl, T. G. Price, M. Schkolnick, P. Griffiths Seli Wade, R. A. Yost

March 1981 **ACM Transactions on Database Systems (TODS)**, Volume 6 Issue 1

Full text available:  pdf(1.57 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

System R supports a high-level relational user language called SQL which may be used by ad hoc users at the terminal or as a data sublanguage in PL/I or COBOL. Host-language programs with embedded SQL statements are processed by a precompiler which replaces the SQL statements by calls to a machine-language access module. The precompilation approach involves work of parsing, name binding, and access path selection from the path of a running program, enabling high performance.

Keywords: compilation, performance measurements, query languages, relational database systems, transactions

36 Outerjoin optimization in multidatabase systems

Arbee L. P. Chen

July 1990

Proceedings of the second international symposium on Databases in parallel and distributed environments

Full text available:  pdf(802.60 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Outerjoin is used in distributed relational multidatabase systems for integrating local schemas to a global schema. The global schema need to be modified, optimized, and decomposed into subqueries at local sites for processing. To integrate local relations in different databases to form a global relation, it is expensive to process. In this paper, based on the definition of the schemas, queries with outerjoin, join, select and project ...

37 Hypermedia Systems and Data Models: Uniform comparison of data models using containment models

E. James Whitehead

June 2002

Proceedings of the thirteenth ACM conference on Hypertext and hypermedia

Full text available:  pdf(300.56 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Containment data models are a subset of entity relationship models in which the allowed relationships are either storage, or inheritance. This paper describes containment relationships, and containment data models, applying a range of monolithic, link server, and hyperbase systems, as well as the Dexter reference model, and the WWW. A key quality of containment data models is their ability to model systems uniformly, allowing a broad ...

Keywords: containment data modeling, hypertext data models

38 Primitives for the manipulation of general subdivisions and the computation of Voronoi

Leonidas Guibas, Jorge Stolfi

April 1985

ACM Transactions on Graphics (TOG), Volume 4 Issue 2

Full text available:  pdf(3.55 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The following problem is discussed: given n points in the plane (the sites) and an arbitrary query point q , find the closest point to q . This problem can be solved by constructing the Voronoi diagram of the given sites and then locating the query point. Two algorithms are given, one that constructs the Voronoi diagram in $O(n \log n)$ time, and another that inserts a new site into an existing Voronoi diagram in $O(\log n)$ time.

Keywords: Euler operators, Voronoi and Delaunay diagrams, closest point, computational topology, convex hulls, nearest neighbours, planar graphs, point location, representation of polyhedra, triangulations

39 DB-1 (databases): data integration: Organizing structured web sources by query schemas: a clustering

Bin He, Tao Tao, Kevin Chen-Chuan Chang

November 2004

Proceedings of the Thirteenth ACM conference on Information and knowledge management

Full text available:  pdf(323.72 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In the recent years, the Web has been rapidly "deepened" with the prevalence of databases online. On this deep Web, structured data is organized by providing structured query interfaces and results. Organizing such structured sources is one of the critical steps toward the integration of heterogeneous Web sources. We observe that, for structured Web sources, attributes in query interfaces are discriminative representative ...

Keywords: data integration, deep Web, hierarchical agglomerative clustering

40 Increasing the instruction fetch rate via multiple branch prediction and a branch address cache

Tse-Yu Yeh, Deborah T. Marr, Yale N. Patt

August 1993

Proceedings of the 7th international conference on Supercomputing

Full text available:  pdf(1.13 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

in a Web-oriented environment*Diana, L.; Ferrari, E.; Hans, G.;*

Web Delivering of Music, 2001. Proceedings. First International Conference on , 23-24 Nov. 2001

Pages:52 - 59

[\[Abstract\]](#) [\[PDF Full-Text \(434 KB\)\]](#) [IEEE CNF](#)**6 Optimized seamless integration of biomolecular data***Eckman, B.A.; Lacroix, Z.; Raschid, L.;*

Bioinformatics and Bioengineering Conference, 2001. Proceedings of the IEEE 2nd International Symposium on , 4-6 Nov. 2001

Pages:23 - 32

[\[Abstract\]](#) [\[PDF Full-Text \(75 KB\)\]](#) [IEEE CNF](#)**7 Improving response time by search pruning in a content-based image retrieval system, using inverted file techniques***Squire, McG.D.; Muller, H.; Muller, W.;*

Content-Based Access of Image and Video Libraries, 1999. (CBAIVL '99) Proceedings. IEEE Workshop on , 22 June 1999

Pages:45 - 49

[\[Abstract\]](#) [\[PDF Full-Text \(104 KB\)\]](#) [IEEE CNF](#)**8 An approach for generating file interfaces***Ebert, A.; Hohenstein, U.; Hoding, M.;*

Database Systems for Advanced Applications, 1999. Proceedings., 6th International Conference on , 19-21 April 1999

Pages:61 - 68

[\[Abstract\]](#) [\[PDF Full-Text \(168 KB\)\]](#) [IEEE CNF](#)**9 Managing dynamic medical data in a distributed mode***Bingyi Hu; Jing Bai; Datian Ye;*

Engineering in Medicine and Biology Society, 1998. Proceedings of the 20th Annual International Conference of the IEEE , Volume: 3 , 29 Oct.-1 Nov. 1998

Pages:1292 - 1294 vol.3

[\[Abstract\]](#) [\[PDF Full-Text \(244 KB\)\]](#) [IEEE CNF](#)**10 Content routing in a network of WAIS servers***Duda, A.; Sheldon, M.A.;*

Distributed Computing Systems, 1994., Proceedings of the 14th International Conference on , 21-24 June 1994

Pages:124 - 132

[\[Abstract\]](#) [\[PDF Full-Text \(716 KB\)\]](#) [IEEE CNF](#)**11 A natural language system and its users***West, V.;*

Natural Language Understanding, IEE Colloquium on , 14 Nov 1988

Pages:6/1 - 6/3

[\[Abstract\]](#) [\[PDF Full-Text \(80 KB\)\]](#) [IEEE CNF](#)

